

# On the use of Markov chain Monte Carlo methods for the sampling of mixture models

## A statistical perspective

Randal Douc · Florian Maire · Jimmy Olsson

Received: date / Accepted: date

**Abstract** In this paper we study asymptotic properties of different data-augmentation-type Markov chain Monte Carlo algorithms sampling from mixture models comprising discrete as well as continuous random variables. Of particular interest to us is the situation where sampling from the conditional distribution of the continuous component given the discrete component is infeasible. In this context, we cast *Carlin & Chib's pseudo-prior method* into the framework of mixture models and discuss and compare different variants of this scheme. We propose a novel algorithm, the *FCC sampler*, which is less computationally demanding than any Metropolised Carlin & Chib-type algorithm. The significant gain of computational efficiency is however obtained at the cost of some asymptotic variance. The performance of the algorithm vis-à-vis alternative schemes is investigated theoretically, using some recent results obtained in [3] for inhomogeneous Markov chains evolving alternately according to two different  $\pi^*$ -reversible Markov transition kernels, as well as numerically.

**Keywords** Asymptotic variance · Carlin & Chib's pseudo-prior method · inhomogeneous Markov chains · Metropolisation · mixture models · Peskun ordering

This work is supported by the Swedish Research Council, Grant 2011-5577.

R. Douc  
Institut Télécom/Télécom SudParis, CNRS UMR 5157  
SAMOVAR, Evry, France  
E-mail: randal.douc@telecom-sudparis.eu

F. Maire  
Institut Télécom/Télécom SudParis, CNRS UMR 5157  
SAMOVAR, Evry, France  
E-mail: florian.maire@telecom-sudparis.eu

J. Olsson  
KTH Royal Institute of Technology, Stockholm, Sweden  
E-mail: jimmyol@kth.se

## 1 Introduction

To sample from *mixture models*  $\pi^*$  comprising a discrete and a continuous random variable, denoted by  $M$  and  $Z$ , respectively, is a fundamental problem in statistics. In this paper we study the use of different data-augmentation-type Markov chain Monte Carlo (MCMC) algorithms for this purpose. Of particular interest to us is the situation where sampling from the conditional distribution of  $Z$  given  $M$  is infeasible.

In many applications the most natural approach to sampling from  $\pi^*$  goes via the *Gibbs sampler*, which samples alternately from the conditional distributions  $M | Z$  and  $Z | M$ . Since the component  $M$  is discrete, the former sampling step is most often feasible (at least when  $\pi^*$  is known up to a normalising constant). On the contrary, drawing  $Z | M$  is in general infeasible; in that case this sampling step is typically *metropolised* by replacing, with a Metropolis-Hastings probability, the value of  $Z$  obtained at the previous iteration by a candidate drawn from some proposal kernel. This yields a so-called *Metropolis-within-Gibbs*—or *hybrid*—sampler.

However, when the modes of the mixture distribution are well-separated, implying a strong correlation between  $M$  and  $Z$ , the Gibbs sampler has in general very limited capacity to move flexibly between the different modes, and exhibits for this reason most often very poor mixing (see [2] for some discussion). Since this problem is due to model dependence, it affects the standard Gibbs as well as the hybrid sampler. In order to cope with this well-known problem, we cast *Carlin & Chib's pseudo-prior method* [1] into the framework of mixture models. The method extends the target model with a set of auxiliary variables that are used for moving the discrete component. When the distribution (deter-

mined by a set of *pseudo-priors*) of the auxiliary variables is chosen optimally (an idealised situation however), the method produces indeed i.i.d. samples from the marginal distribution of  $M$  under  $\pi^*$ . Given  $M$ , the  $Z$  component is sampled from  $Z \mid M$  in accordance with the Gibbs sampler, with possible metropolisation in the case where exact sampling is infeasible. The latter scheme will be referred as the *Metropolised Carlin & Chib-type (MCC) sampler*.

Surprisingly, it turns out that passing directly and deterministically the value of the  $M$ th auxiliary variable, obtained through sampling from the pseudo-priors at the beginning of the loop, to the  $Z$  component yields a Markov chain that is still  $\pi^*$ -reversible (see Theorem 4), and using some novel results obtained in [3] on the comparison of asymptotic variance for inhomogeneous Markov chains we are able to prove (see Theorem 5) that this novel MCMC algorithm, referred to as the *Frozen Carlin & Chib-type (FCC) sampler*, generates a Markov chain whose sample path averages have always higher asymptotic variance than those of the MCC sampler for a large class of objective functions. This is well in line with our expectations, as the MCC sampler “refreshes” more often the  $Z$  component. On the other hand, since this component is already modified through sampling from the pseudo-priors, which, when well-designed, should be close to the true conditional distribution  $Z \mid M$ , we may expect that the additional mixing provided by the MCC sampler is only marginal. This is also confirmed by our simulations, which indicate only a small advantage of the MCC sampler to the FCC sampler in terms of autocorrelation. As the FCC algorithm omits completely the Metropolis-Hastings operation of the MCC sampler, it is considerably more computationally efficient. Thus, we consider the FCC sampler as a strong alternative to the MCC sampler in terms of efficiency (variance per unit CPU).

The paper is structured as follows: in Section 2 we introduce some notation and describe the mixture model framework under consideration. Section 3 describes the Carlin & Chib-type MCMC samplers studied in the paper. In Section 4 we prove that the involved algorithms are indeed  $\pi^*$ -reversible and provide a theoretical comparison of the MCC and FCC samplers. Finally, in the implementation part, Section 5, we illustrate and compare numerically the algorithms on a two examples: a mixture of Gaussian distributions and a model where the mixture variables are only partially observed.

## 2 Preliminaries

### 2.1 Notation

We assume throughout the paper that all variables are defined on a common probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ . We will use upper case for random variables and lower case for realisations of the same, and write “ $X \rightsquigarrow x$ ” when  $x$  is realisation of  $X$ . We write “ $X \sim \mu$ ” to indicate that the random variable  $X$  is distributed according to the probability measure  $\mu$ . For any  $\mu$ -integrable function  $h$  we let  $\mu(h) := \int h(x)\mu(dx)$  be the expectation of  $h(X)$  under  $\mu$ . Similarly, for Markov transition kernels  $M$  we write  $Mf(x) := \int f(x')M(x, dx')$  whenever this integral is well-defined. For any two probability measures  $\mu$  and  $\mu'$  defined on some measurable spaces  $(X, \mathcal{X})$  and  $(X', \mathcal{X}')$ , respectively, we denote by  $\mu(dx)\mu'(dx')$  the product measure  $\mu \otimes \mu'(dx \times dx')$  on  $(X \times X', \mathcal{X} \otimes \mathcal{X}')$ . For  $(m, n) \in \mathbb{Z}^2$  such that  $m \leq n$ , we denote by  $\llbracket m, n \rrbracket := \{m, m+1, \dots, n\} \subset \mathbb{Z}$ . Moreover, we denote by  $\mathbb{N}^* := \mathbb{N} \setminus \{0\}$  the set of positive integers.

Finally, given some probability measure  $\pi$  on  $(X, \mathcal{X})$  we recall, firstly, that a Markov transition kernel  $M$  is called  $\pi$ -reversible if  $\pi(dx)M(x, dx') = \pi(x')M(x', dx)$  and, secondly, that  $\pi$ -reversibility of  $M$  implies straightforwardly that this kernel allows  $\pi$  as a stationary distribution.

### 2.2 Mixture models

Throughout this paper, our main objective is to sample a probability distribution  $\pi^*$  on some product space  $Y := \llbracket 1, n \rrbracket \times Z$ , where  $(Z, \mathcal{Z})$  is some (typically uncountable) measurable space, associated with the  $\sigma$ -field  $\mathcal{Y} := 2^{\llbracket 1, n \rrbracket} \otimes \mathcal{Z}$ . Thus, a  $\pi^*$ -distributed random variable  $Y = (M, Z)$  comprises an  $\llbracket 1, n \rrbracket$ -valued (discrete) random variable  $M$  and a  $Z$ -valued (typically continuous) random variable  $Z$ .

In the following we assume that  $\pi^*(dm \times dz)$  is dominated by a product measure  $|dm|\nu(dz)$ , where  $|dm|$  denotes the counting measure on  $\llbracket 1, n \rrbracket$  and  $\nu$  is some nonnegative measure on  $(Z, \mathcal{Z})$ , and denote by  $\pi^*(m, z)$  the corresponding density function on  $\llbracket 1, n \rrbracket \times Z$ . We may then define the conditional density and probability functions

$$\begin{aligned} \pi^*(m \mid z) &:= \frac{\pi^*(m, z)}{\sum_{m'=1}^n \pi^*(m', z)}, \\ \pi^*(z \mid m) &:= \frac{\pi^*(m, z)}{\int \pi^*(m, z')\nu(dz')} \end{aligned} \tag{1}$$

(w.r.t.  $|dm|$  and  $\nu$ , respectively) on  $\llbracket 1, n \rrbracket$  and  $Z$ , respectively. We also define the marginal probability function

$$\pi^*(m) := \int \pi^*(m, z) \nu(dz)$$

(w.r.t.  $|dm|$ ) on  $\llbracket 1, n \rrbracket$ .

### 3 Markov chain Monte Carlo methods for mixture models

Using the conditional distributions (1), a natural way of sampling  $\pi^*$  consists in implementing a standard Gibbs sampler simulating a Markov chain  $\{Y_k^{(1)}; k \in \mathbb{N}\}$  with transitions described by the following algorithm.

---

#### Algorithm 1 Gibbs sampler

---

**Require:**  $Y_k^{(1)} = (m, z)$ ,  
 (i) draw  $M' \sim \pi^*(dm | z)$  and call the outcome  $m'$  (abbr.  $\rightsquigarrow m'$ ),  
 (ii) draw  $Z' \sim \pi^*(dz | m') \rightsquigarrow z'$ ,  
 (iii) set  $Y_{k+1}^{(1)} \leftarrow (m', z')$ .

---

**Remark 1** *Since  $M$  is a discrete random variable it is most often possible to sample  $M \sim \pi^*(dm | z)$ . In contrast, sampling  $Z \sim \pi^*(dz | m)$  is not always possible. In that case, one may replace Step (ii) by a Metropolis-Hastings step, yielding a Metropolis-within-Gibbs algorithm (see [6, section 10.3.3] for details).*

Using the output Algorithm 1, any expectation  $\pi^*(f)$ , where  $f$  is some  $\pi^*$ -integrable objective function on  $Y$ , can be estimated by the sample path average

$$\hat{\pi}_n^{(1)}(f) := \frac{1}{n} \sum_{k=0}^{n-1} f(Y_k^{(1)}).$$

Even though Algorithm 1 generates a Markov chain  $\{Y_k^{(1)}; k \in \mathbb{N}\}$  with stationary distribution  $\pi^*$ , the discrete component  $\{M_k^{(1)}; k \in \mathbb{N}\}$  tends in practice to get stuck in a few states. Indeed, when the variable  $Z$  is sampled from its conditional distribution given  $M = m$ , the probability of jumping to another index  $m' \neq m$  is proportional to  $\pi^*(m', z)$ , which may be very low when the index component  $M$  is informative concerning the localisation of  $Z$ . This will lead to poor mixing and, consequently, high variance of any estimator  $\hat{\pi}_n^{(1)}(f)$ . To be specific, let  $h$  be some function on  $\llbracket 1, n \rrbracket$  and assume that we run the Gibbs sampler in Algorithm 1 to estimate  $\int h(m) \pi^*(dm)$ . Then  $\{M_k^{(1)}; k \in \mathbb{N}\}$  is itself a Markov chain with transition kernel  $G$ , say, and,

starting with  $M_0^{(1)} \sim \pi^*(dm)$ , we have

$$\begin{aligned} \text{Cov}(h(M_0^{(1)}), h(M_1^{(1)})) &= \text{Cov}(h(M_0^{(1)}), Gh(M_0^{(1)})) \\ &= \int \pi^*(dz) \left( \int \pi^*(dm | z) h(m) \right)^2 \geq 0. \end{aligned}$$

Combining this with the fact that  $G$  is  $\pi^*(dm)$ -reversible, we obtain

$$\begin{aligned} \text{Cov}(h(M_0^{(1)}), h(M_{2k+1}^{(1)})) \\ = \text{Cov}(G^k h(M_0^{(1)}), G(G^k h)(M_0^{(1)})) \geq 0. \end{aligned}$$

Moreover, using again that  $G$  is  $\pi^*(dm)$ -reversible,

$$\begin{aligned} \text{Cov}(h(M_0^{(1)}), h(M_{2k}^{(1)})) \\ = \text{Cov}(G^k h(M_0^{(1)}), G^k h(M_0^{(1)})) \geq 0. \end{aligned}$$

Finally, letting  $f(m, z) \equiv h(m)$ , we obtain

$$\text{Var} \left( \sqrt{n} \hat{\pi}_n^{(1)}(f) \right) \geq \text{Var} \left( h(M_0^{(1)}) \right),$$

showing that the Gibbs sampler approximates the index less accurately than i.i.d. sampling from  $\pi^*(dm)$ .

The *pseudo-prior method* of B. P. Carlin and S. Chib [1] was introduced in the context of model selection and can successfully be adapted to mixture models. By introducing some auxiliary variables, this method increases the number of moves of the index component. The algorithm may be regarded as a Gibbs sampler-based data-augmentation algorithm targeting the distribution  $\pi$  defined on the extended state space  $\llbracket 1, n \rrbracket \times Z^n$  by

$$\pi(dm \times du) := \pi^*(dm \times du_m) \bigotimes_{j \neq m} \rho_j(du_j), \quad (2)$$

where  $u = (u_1, \dots, u_n) \in Z^n$  and the probability measures  $\{\rho_j; j \in \llbracket 1, n \rrbracket\}$  are referred to as *pseudo-priors* or *linking densities* (the terminology comes from [1]). We assume that also the pseudo-priors are dominated jointly by the nonnegative measure  $\nu$  and use the same symbols  $\rho_j$  for denoting the corresponding densities. As a consequence, also the measure  $\pi$  defined in (2) has a density

$$\pi(m, u) := \pi^*(m, u_m) \prod_{j \neq m} \rho_j(u_j) \quad ((m, u) \in \llbracket 1, n \rrbracket \times Z^n)$$

with respect to  $|dm| \otimes \nu^{\otimes n}(du)$ .

The choice of the pseudo-priors is tuned by the user provided that these are analytically tractable and can be sampled from. Denote by  $\{Y_k^{(2)}; k \in \mathbb{N}\}$  the Markov chain generated by this algorithm, which we will in the following refer to as the *Carlin & Chib-type* (CC-type)

sampler and whose transitions comprise the  $n + 1$  sub-steps described in Algorithm 2 below.

---

**Algorithm 2** CC-type sampler
 

---

**Require:**  $Y_k^{(2)} = (m, u_m)$ ,  
 (i) for all  $j \neq m$ , draw  $U_j \sim \rho_j \rightsquigarrow u_j$ ,  
 (ii) draw  $M' \sim \pi(dm | u) \rightsquigarrow m'$ ,  
 (iii) draw  $U'_{m'} \sim \pi^*(dz | m') \rightsquigarrow u'_{m'}$ ,  
 (iv) set  $Y_{k+1}^{(2)} \leftarrow (m', u'_{m'})$ .

---

Intuitively, Algorithm 2 allows different models to be visited more frequently than in the Gibbs sampler; indeed, in Step (ii) the probability of moving to the index  $m'$  is

$$\pi(m' | u) \propto \pi^*(m', u_{m'}) / \rho_{m'}(u_{m'}), \quad (3)$$

where the right hand side is close to  $\pi^*(m')$  if the pseudo-priors are chosen such that  $\rho_\ell(z)$  is close to  $\pi^*(z | \ell)$  for all  $(\ell, z) \in \llbracket 1, n \rrbracket \times \mathcal{Z}$ . The optimal case where  $\rho_\ell(z) \equiv \pi^*(z | \ell)$  implies, via (3), that

$$\pi(m' | u) \propto \pi^*(m').$$

Thus, in this case, Step (ii) draws actually  $M'$  according to the *exact* marginal  $\pi^*(dm)$  of the class index random variable regardless the value of  $u$ , which implies that the algorithm simulates *i.i.d. samples* according to  $\pi^*$ . This actually gives a more efficient approximation than that produced by the Gibbs sampler, whose variance w.r.t. the index component is, as we remarked previously, always larger than that obtained through i.i.d. sampling from  $\pi^*(dm)$ . However, this ideal situation requires the quantity  $\pi^*(z | \ell)$  to be tractable, which is typically not the case.

As in Remark 1, one may replace Step (iii) in Algorithm 2 by a Metropolis-Hastings step if sampling from  $\pi^*(dz | m')$  is infeasible. This is most often the case when  $\pi^*(dm \times dz)$  is the *a posteriori* distribution of  $(M, Z)$  conditionally on one or several observations (see Section 5.2 for an example). In that case  $\pi^*$  is known only up to a normalizing constant, which prevents sampling from the conditional density  $\pi^*(dz | m')$ . The resulting algorithm will in the following be referred to as the Metropolised CC-type (MCC) sampler and is presented in Algorithm 3, where  $\{R_\ell; \ell \in \llbracket 1, n \rrbracket\}$  is set of proposal kernels on  $\mathcal{Z} \times \mathcal{Z}$ . Assume for simplicity that all these kernels are jointly dominated by the reference measure  $\nu$  and denote by  $\{r_\ell; \ell \in \llbracket 1, n \rrbracket\}$  the corresponding transition densities with respect to this measure. Introducing also the Metropolis-Hastings acceptance probability

$$\alpha_\ell(u, z) := \frac{\pi^*(\ell, z)r_\ell(z, u)}{\pi^*(\ell, u)r_\ell(u, z)} \wedge 1 \quad ((\ell, u, z) \in \llbracket 1, n \rrbracket \times \mathcal{Z}^2),$$

the MCC algorithm is described as follows.

---

**Algorithm 3** MCC sampler
 

---

**Require:**  $Y_k^{(3)} = (m, u_m)$ ,  
 (i) for all  $j \neq m$ , draw  $U_j \sim \rho_j \rightsquigarrow u_j$ ,  
 (ii) draw  $M' \sim \pi(dm | u) \rightsquigarrow m'$ ,  
 (iii.1) draw  $Z \sim R_{m'}(u_{m'}, dz) \rightsquigarrow z$ ,  
 (iii.2) set  $U'_{m'} \leftarrow \begin{cases} z & \text{w. pr. } \alpha_{m'}(u_{m'}, z), \\ u_{m'} & \text{otherwise,} \end{cases} \rightsquigarrow u'_{m'}$ ,  
 (iv) set  $Y_{k+1}^{(3)} \leftarrow (m', u'_{m'})$ .

---

Note that Step (iii) generates, given  $u_{m'}$ ,  $U'_{m'} \sim K_{m'}(u_{m'}, du')$ , where

$$\begin{aligned} K_\ell(u, du') &:= R_\ell(u, du') \alpha_\ell(u, u') \\ &+ \delta_u(du') \left( 1 - \int R_\ell(u, du'') \alpha_\ell(u, u'') \right) \\ &((u, \ell) \in \llbracket 1, n \rrbracket). \end{aligned} \quad (5)$$

It can be easily checked (using (4)) that  $K_{m'}$  is indeed a Metropolis-Hastings kernel with respect to  $\pi(dz | m')$ ; it is thus  $\pi(dz | m')$ -reversible.

Remarkably, it turns out that Step (iii) in Algorithm 3 may be omitted, which may, in some cases, imply a significant gain of computational complexity. The MCC sampler then simplifies to what we will refer to as the Frozen CC-type (FCC) sampler which is described formally as follows.

---

**Algorithm 4** FCC sampler
 

---

**Require:**  $Y_k^{(4)} = (m, u_m)$ ,  
 (i) for all  $j \neq m$ , draw  $U_j \sim \rho_j \rightsquigarrow u_j$ ,  
 (ii) draw  $M' \sim \pi(dm | u) \rightsquigarrow m'$ ,  
 (iii)  $u'_{m'} \leftarrow u_{m'}$ ,  
 (iv) set  $Y_{k+1}^{(4)} \leftarrow (m', u'_{m'})$ .

---

As remarked in Theorem 4 in the next section, this novel algorithm produces a Markov chain  $\{Y_k^{(4)}; k \in \mathbb{N}\}$  that indeed admits  $\pi^*$  as an invariant distribution. Nevertheless, as stated in Theorem 5 below, this algorithm is always less efficient in terms of asymptotic variance than the corresponding MCC sampler. Intuitively, this stems from the fact that once the index  $M'$  is drawn, the associated continuous component is selected deterministically without being “refreshed” (on the contrary to Step (iii) in Algorithm 3). Nevertheless, as mentioned above, Algorithm 4 skips completely the Metropolis-Hastings step (Step (iii)) of Algorithm 3 and is therefore considerably less demanding from a computational

point of view. In addition, as our numerical simulations indicate that the gain of the asymptotic variance obtained by refreshing, as in the MCC sampler, this component instead of freezing the same as in the FCC sampler seems to be limited (see Section 5 for details), we definitely regard the FCC algorithm as a strong challenger of the MCC sampler.

## 4 Theoretical results

### 4.1 Comparison of asymptotic variance of inhomogeneous Markov chains

In this section we recall briefly the main result of [3, Theorem 4], which is propelling the coming analysis. The following—now classical—orderings of Markov kernels turns out to be highly useful.

**Definition 1** Let  $P_0$  and  $P_1$  be Markov transition kernels on some state space  $(X, \mathcal{X})$  with common invariant distribution  $\pi$ . We say that  $P_1$  *dominates*  $P_0$

- on the *off-diagonal*, denoted  $P_1 \succeq P_0$ , if for all  $A \in \mathcal{X}$  and  $\pi$ -a.s. all  $x \in X$ ,

$$P_1(x, A \setminus \{x\}) \geq P_0(x, A \setminus \{x\}).$$

- in the *covariance ordering*, denoted  $P_1 \succcurlyeq P_0$ , if for all  $f \in L^2(\pi)$ ,

$$\int f(x) P_1 f(x) \pi(dx) \leq \int f(x) P_0 f(x) \pi(dx).$$

The covariance ordering, which was introduced implicitly in [7, p. 5] and formalised in [4], is an extension of the off-diagonal ordering, since, according to [7, Lemma 3],  $P_1 \succeq P_0$  implies  $P_1 \succcurlyeq P_0$ . Moreover, it turns out that for reversible kernels,  $P_1 \succcurlyeq P_0$  implies that the asymptotic variance of sample path averages of chains generated by  $P_1$  is smaller than or equal to that of chains generated by  $P_0$  (see the proof of [7, Theorem 4]).

In algorithms of Gibbs-type, the ordering in Definition 1 is usually not applicable, since the fact that all candidates are accepted with probability one prevents the chain from remaining in the same state. The ordering is however still meaningful when a component is discrete.

In the following, let  $P_i$  and  $Q_i$ ,  $i \in \llbracket 0, 1 \rrbracket$ , be Markov transition kernels on  $(X, \mathcal{X})$  and let  $\{X_k^{(0)}; k \in \mathbb{N}\}$  and  $\{X_k^{(1)}; k \in \mathbb{N}\}$  be inhomogeneous Markov chains evolving as follows:

$$X_0^{(i)} \xrightarrow{P_i} X_1^{(i)} \xrightarrow{Q_i} X_2^{(i)} \xrightarrow{P_i} X_3^{(i)} \xrightarrow{Q_i} \dots \quad (6)$$

This means that for all  $k \in \mathbb{N}$  and  $i \in \{0, 1\}$ ,

$$\begin{aligned} - \mathbb{P} \left( X_{2k+1}^{(i)} \in dx \mid \mathcal{F}_{2k}^{(i)} \right) &= P_i(X_{2k}^{(i)}, dx), \\ - \mathbb{P} \left( X_{2k+2}^{(i)} \in dx \mid \mathcal{F}_{2k+1}^{(i)} \right) &= Q_i(X_{2k+1}^{(i)}, dx), \end{aligned}$$

where  $\mathcal{F}_n^{(i)} := \sigma(X_0^{(i)}, \dots, X_n^{(i)})$ ,  $n \in \mathbb{N}$ . Now, impose the following assumption.

- (A1)** (i)  $P_i$  and  $Q_i$ ,  $i \in \llbracket 0, 1 \rrbracket$ , are  $\pi$ -reversible,  
(ii)  $P_1 \succcurlyeq P_0$  and  $Q_1 \succcurlyeq Q_0$ .

As mentioned above,  $P_1 \succeq P_0$  implies  $P_1 \succcurlyeq P_0$ ; thus, in practice, a sufficient condition for **(A1)**(ii) is that  $P_1 \succeq P_0$  and  $Q_1 \succeq Q_0$ . Under these assumptions, [3] established the following result.

**Theorem 2 ([3])** Assume that  $P_i$  and  $Q_i$ ,  $i \in \llbracket 0, 1 \rrbracket$ , satisfy **(A1)** and let  $\{X_k^{(i)}; k \in \mathbb{N}\}$ ,  $i \in \llbracket 0, 1 \rrbracket$ , be Markov chains evolving as in (6) with initial distribution  $\pi$ . Then for all  $f \in L^2(\pi)$  such that for  $i \in \llbracket 0, 1 \rrbracket$ ,

$$\sum_{k=1}^{\infty} \left( |\text{Cov}(f(X_0^{(i)}), f(X_k^{(i)}))| + |\text{Cov}(f(X_1^{(i)}), f(X_{k+1}^{(i)}))| \right) < \infty \quad (7)$$

it holds that

$$v_1(f) \leq v_0(f), \quad (8)$$

where

$$v_i(f) := \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left( \sum_{k=0}^{n-1} f(X_k^{(i)}) \right) \quad (i \in \llbracket 0, 1 \rrbracket). \quad (9)$$

**Remark 3** As shown in [3, Proposition 9], under the assumption that the product kernels  $P_i Q_i$ ,  $i \in \llbracket 0, 1 \rrbracket$ , are both  $V$ -geometrically ergodic (according to Definition 7 in the same paper), the absolute summability assumption (7) holds true for all objective functions  $f$  such that  $f$  and  $P_i f$ ,  $i \in \llbracket 0, 1 \rrbracket$ , have all bounded  $\sqrt{V}$ -norm; see again [3] for details.

### 4.2 The MCC sampler vs. the FCC sampler

In the light of the remarks following Algorithm 4 it is reasonable to assume that the CC-type sampler (Algorithm 2) and the MCC sampler provides more accurate estimates than the FCC sampler. However, since neither  $\{Y_k^{(3)}; k \in \mathbb{N}\}$  nor  $\{Y_k^{(4)}; k \in \mathbb{N}\}$  are  $\pi^*$ -reversible, [7, Theorem 4] does not allow these two algorithms to be compared. Nevertheless, using Theorem 2 we may provide a theoretical justification advocating the MCC and CC-type samplers ahead of the FCC sampler in terms of asymptotic variance. To do this we first embed  $\{Y_k^{(3)}; k \in \mathbb{N}\}$  and  $\{Y_k^{(4)}; k \in \mathbb{N}\}$  into inhomogeneous  $\pi^*$ -reversible Markov chains  $\{X_k^{(3)}; k \in \mathbb{N}\}$  and

$\{X_k^{(4)}; k \in \mathbb{N}\}$  defined on  $\mathcal{Y} = \llbracket 1, n \rrbracket \times \mathcal{Z}$  through, for  $i \in \llbracket 3, 4 \rrbracket$ :

$$\begin{aligned} X_{2k}^{(i)} &= \begin{pmatrix} M_k^{(i)} \\ Z_k^{(i)} \end{pmatrix} \xrightarrow{P_i} X_{2k+1}^{(i)} = \begin{pmatrix} \check{M}_{k+1}^{(i)} \\ \check{Z}_{k+1}^{(i)} \end{pmatrix} \\ &\xrightarrow{Q_i} X_{2k+2}^{(i)} = \begin{pmatrix} M_{k+1}^{(i)} \\ Z_{k+1}^{(i)} \end{pmatrix} \xrightarrow{P_i} \dots \quad (10) \end{aligned}$$

Here we have defined the kernels

- $P_3((m, z), d\check{m} \times d\check{z})$   
 $:= \int \cdots \int \left( \prod_{j \neq m} \rho_j(du_j) \right) \delta_z(du_m) \pi(d\check{m} | u) \delta_{u_{\check{m}}}(d\check{z}),$
- $P_4 := P_3,$
- $Q_3((\check{m}, \check{z}), dm \times dz) := \delta_{\check{m}}(dm) K_{\check{m}}(\check{z}, dz)$  (where  $K_{\check{m}}$  is defined in (5)),
- $Q_4((\check{m}, \check{z}), dm \times dz) := \delta_{\check{m}}(dm) \delta_{\check{z}}(dz).$

Setting  $Y_k^{(i)} := (M_k^{(i)}, Z_k^{(i)})$ ,  $k \in \mathbb{N}$ ,  $i \in \llbracket 1, 2 \rrbracket$ , it can be checked easily that  $\{Y_k^{(3)}; k \in \mathbb{N}\}$  and  $\{Y_k^{(4)}; k \in \mathbb{N}\}$  have indeed exactly the same distribution as the output of Algorithm 3 and Algorithm 4, respectively.

**Theorem 4** *The Markov chains  $\{Y_k^{(3)}; k \in \mathbb{N}\}$  and  $\{Y_k^{(4)}; k \in \mathbb{N}\}$  generated by Algorithm 3 and Algorithm 4, respectively, have  $\pi^*$  as invariant distribution.*

*Proof.* The result is established by noting that  $Q_4$  defined above is reversible with respect to any distribution, and in particular it is  $\pi^*$ -reversible. Moreover, according to Lemma 6 (below),  $P_3 = P_4$  and  $Q_3$  are also  $\pi^*$ -reversible. The statement of the theorem follows.  $\square$

**Theorem 5** *Let  $\{X_k^{(i)}; k \in \mathbb{N}\}$ ,  $i \in \llbracket 3, 4 \rrbracket$ , be the Markov chains (10) starting with  $X_0^{(i)} \sim \pi^*$  for  $i \in \llbracket 3, 4 \rrbracket$ . Then for all real-valued functions  $h$  such that for  $i \in \llbracket 3, 4 \rrbracket$ ,*

$$\sum_{k=1}^{\infty} |\text{Cov}(h(M_0^{(i)}), h(M_k^{(i)}))| < \infty$$

*it holds that*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left( \sum_{k=1}^n h(M_k^{(3)}) \right) \leq \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left( \sum_{k=1}^n h(M_k^{(4)}) \right).$$

*Proof.* By Theorem 4, the processes  $\{X_k^{(1)}; k \in \mathbb{N}\}$  and  $\{X_k^{(2)}; k \in \mathbb{N}\}$  are both inhomogeneous Markov chains that evolve alternately according to the  $\pi^*$ -reversible kernels  $P_i$  and  $Q_i$ ,  $i \in \llbracket 1, 2 \rrbracket$ . Moreover,  $P_3 = P_4 \succeq P_4$ , and since  $Q_4$  has no off-diagonal component,  $Q_3 \succeq Q_4$ . Now, define Markov chains  $\{X_k^{(i)}; k \in \mathbb{N}\}$ ,  $i \in \llbracket 3, 4 \rrbracket$ , as in (10) with  $X_0^{(i)} \sim \pi$ , and set  $f(m, z) \equiv$

$h(m)$ . By construction,  $\check{M}_k^{(i)} = M_k^{(i)}$  for  $i \in \llbracket 3, 4 \rrbracket$  and  $k \in \mathbb{N}^*$ , implying that

$$\begin{aligned} &\sum_{k=1}^{\infty} (|\text{Cov}(f(X_0^{(i)}), f(X_k^{(i)}))| + |\text{Cov}(f(X_1^{(i)}), f(X_{k+1}^{(i)}))|) \\ &= \pi f^2 - \pi^2 f + 4 \sum_{k=1}^{\infty} |\text{Cov}(h(M_0^{(i)}), h(M_k^{(i)}))| < \infty. \end{aligned} \quad (11)$$

Moreover, for all  $n \in \mathbb{N}^*$  and  $i \in \llbracket 3, 4 \rrbracket$ ,

$$\begin{aligned} \text{Var} \left( \sum_{k=1}^n h(M_k^{(i)}) \right) &= \text{Var} \left( \sum_{k=1}^n h(\check{M}_k^{(i)}) \right) \\ &= \frac{1}{4} \text{Var} \left( \sum_{k=1}^{2n} f(X_k^{(i)}) \right), \end{aligned}$$

which implies, by (11), that for  $i \in \llbracket 3, 4 \rrbracket$ ,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left( \sum_{k=1}^n h(M_k^{(i)}) \right) = \frac{1}{2} \lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left( \sum_{k=1}^n f(X_k^{(i)}) \right).$$

Finally, by (11) we may apply Theorem 2 to the chains  $\{X_k^{(i)}; k \in \mathbb{N}\}$ ,  $i \in \llbracket 0, 1 \rrbracket$ , which establishes immediately the statement of the theorem.  $\square$

**Lemma 6** *The Markov kernels  $P_3$  and  $Q_3$  are both  $\pi^*$ -reversible.*

*Proof.* Write, using the identity

$$\begin{aligned} \nu(dz) \delta_z(du_m) \delta_{u_{\check{m}}}(d\check{z}) \prod_{j \neq m} \nu(du_j) \\ = \delta_{u_m}(dz) \delta_{u_{\check{m}}}(d\check{z}) \prod_{j=1}^n \nu(du_j), \end{aligned}$$

for any nonnegative measurable function  $f$  on  $(\mathcal{Y}, \mathcal{Y})$ ,

$$\begin{aligned} &\iint f(y, \check{y}) \pi^*(dy) P_3(y, d\check{y}) \\ &= \int \cdots \int f(y, \check{y}) \pi^*(m, z) |dm| \nu(dz) \delta_z(du_m) \\ &\quad \times \left( \prod_{j \neq m} \rho_j(du_j) \right) \frac{\pi^*(\check{m}, u_{\check{m}}) \prod_{j \neq \check{m}} \rho_j(u_j)}{\sum_{k=1}^n \pi^*(k, u_k) \prod_{\ell \neq k} \rho_\ell(u_\ell)} \\ &\quad \times |d\check{m}| \delta_{u_{\check{m}}}(d\check{z}) \\ &= \int \cdots \int f(y, \check{y}) \pi^*(m, z) \pi^*(\check{m}, u_{\check{m}}) \\ &\quad \times \frac{\prod_{j \neq m} \rho_j(u_j) \prod_{j \neq \check{m}} \rho_j(u_j)}{\sum_{k=1}^n \pi^*(k, u_k) \prod_{\ell \neq k} \rho_\ell(u_\ell)} \\ &\quad \times \left( \prod_{j=1}^n \nu(du_j) \right) |dm| |d\check{m}| \delta_{u_m}(dz) \delta_{u_{\check{m}}}(d\check{z}). \end{aligned}$$

Thus, integrating first over  $z$  and  $\tilde{z}$  and defining

$$A(m, \tilde{m}, u) := \pi^*(m, u_m) \pi^*(\tilde{m}, u_{\tilde{m}}) \frac{\prod_{j \neq m} \rho_j(u_j) \times \prod_{j \neq \tilde{m}} \rho_j(u_j)}{\sum_{k=1}^n \pi^*(k, u_k) \prod_{\ell \neq k} \rho_\ell(u_\ell)}$$

yields

$$\begin{aligned} & \iint f(y, \tilde{y}) \pi^*(dy) P_2(y, d\tilde{y}) \\ &= \int \cdots \int f((m, u_m), (\tilde{m}, u_{\tilde{m}})) A(m, \tilde{m}, u) \\ & \quad \times \prod_{j=1}^n \nu(du_j) |dm| |d\tilde{m}|. \end{aligned}$$

Now, the symmetry  $A(m, \tilde{m}, u) = A(\tilde{m}, m, u)$  implies the identity

$$\iint f(y, \tilde{y}) \pi^*(dy) P_3(y, d\tilde{y}) = \iint f(y, \tilde{y}) \pi^*(d\tilde{y}) P_3(\tilde{y}, dy), \quad (12)$$

and as  $f$  was chosen arbitrarily, (12) implies that

$$\pi^*(dy) P_3(y, d\tilde{y}) = \pi^*(d\tilde{y}) P_3(\tilde{y}, dy),$$

which establishes the  $\pi^*$ -reversibility of  $P_3$ .

We show that  $Q_3$  is  $\pi^*$ -reversible. Again, let  $f$  be some nonnegative measurable function on  $(Y, \mathcal{Y})$ . Then, using that  $K_m$  is reversible with respect to  $\pi^*(dz | m)$  for all  $m \in \llbracket 1, n \rrbracket$ , we obtain, denoting  $\tilde{y} := (\tilde{m}, \tilde{z})$  and  $y := (m, z)$ ,

$$\begin{aligned} & \int \cdots \int f(\tilde{y}, y) \pi^*(d\tilde{y}) K_{\tilde{m}}(\tilde{z}, dz) \delta_{\tilde{m}}(dm) \\ &= \int \cdots \int f(\tilde{y}, y) \pi^*(d\tilde{m}) \pi^*(d\tilde{z} | \tilde{m}) K_{\tilde{m}}(\tilde{z}, dz) \delta_{\tilde{m}}(dm) \\ &= \int \cdots \int f(\tilde{y}, y) \pi^*(d\tilde{m}) \pi^*(dz | \tilde{m}) K_{\tilde{m}}(z, d\tilde{z}) \delta_{\tilde{m}}(dm) \\ &= \int \cdots \int f(\tilde{y}, y) \pi^*(dm) \pi^*(dz | m) K_m(z, d\tilde{z}) \delta_m(d\tilde{m}). \end{aligned}$$

This implies

$$\begin{aligned} & \iint f(\tilde{y}, y) \pi^*(dy) Q_3(y, d\tilde{y}) \\ &= \iint f(\tilde{y}, y) \pi^*(d\tilde{y}) Q_3(\tilde{y}, dy), \end{aligned}$$

which completes the proof.  $\square$

## 5 Numerical illustrations

In this section we compare numerically the performances of the different algorithms described in the previous section. The comparisons will be based on two different models: firstly, a simple toy model consisting of a mixture of two Gaussian strata and, secondly, a model where only partial observations of the mixture variables are available. All implementations are in MATLAB, running on a MacBook Air with a 1.8 GHz Inter Core i7 processor.

### 5.1 Mixture of Gaussian strata

Let  $Y = \llbracket 1, 2 \rrbracket \times \mathbb{R}$  (i.e.  $Z = \mathbb{R}$  in this case) and consider a pair of random variables  $(M, Z)$  distributed according to the Gaussian mixture model

$$\pi^*(m, z) = \frac{1}{2} \phi(z; \mu_m, \sigma^2) \quad ((m, z) \in Y), \quad (13)$$

where  $\sigma > 0$ ,  $(\mu_1, \mu_2) = (-1, 1)$ , and  $\phi(z; \mu, \sigma^2)$  denotes the Gaussian probability density function with mean  $\mu$  and variance  $\sigma^2$ . Even though it is straightforward to generate i.i.d. samples from this simple toy model, we use it for illustrating and comparing the performances of the algorithms proposed in the previous; in particular, since  $\pi^*(z | m)$  is simply a Gaussian distribution in this case, it is possible to execute Step (iii) in Algorithm 2 (which is, as mentioned, far from always the case; see the next example). For small values of  $\sigma$ , such as the value  $\sigma = \sqrt{2}$  used in this simulation, the two modes are well-separated, implying a strong correlation between the discrete and continuous components. As a consequence we may expect the naive Gibbs sampler to exhibit a very sub-optimal performance in this case. In order to improve mixing we introduced Gaussian pseudo-priors

$$\rho_\ell(u) := \phi(u; \tilde{\mu}_\ell, \tilde{\sigma}_\ell^2) \quad ((\ell, u) \in \llbracket 1, 2 \rrbracket \times \mathbb{R})$$

on  $\mathbb{R}$ , where  $(\tilde{\mu}_1, \tilde{\mu}_2, \tilde{\sigma}_1^2, \tilde{\sigma}_2^2) = (-.5, .5, .15, .25)$ , and executed, using these pseudo-priors, Algorithm 2, Algorithm 3, and Algorithm 4. Moreover, the naive Gibbs sampler was implemented for comparison. Algorithm 3 used the proposal

$$R_\ell(u, dz) = \rho_\ell(dz) \quad ((\ell, u) \in \llbracket 1, 2 \rrbracket \times \mathbb{R}),$$

yielding an algorithm that can be viewed as a hybrid between Algorithm 2 and Algorithm 4 in the sense that it “refreshes” randomly the continuous component  $U_{m'}$  obtained after Step (ii) by replacing, with the Metropolis-Hastings probability  $\alpha_{m'}$ , the same by a draw from  $\rho'_{m'}$ . Cf. Algorithm 2 and Algorithm 4, where  $U_{m'}$  is refreshed systematically according to  $\rho'_{m'}$  and kept frozen,

respectively. For each of these algorithms we generated an MCMC trajectory comprising 101,000 iterations (where the first 1,000 iterations were regarded as burn-in and discarded) and estimated the corresponding autocorrelation functions. The outcome, which is displayed in Figure 1 below, indicates increasing autocorrelation for the CC, MCC, FCC, and Gibbs algorithms, respectively, confirming completely the theoretical results obtained in the previous section. Interestingly, the FCC algorithm has, despite being close to twice as efficient in terms of CPU with our implementation, only slightly higher autocorrelation than the MCC algorithm (the same applies to both the components). As expected, the Gibbs sampler suffers from very large autocorrelation as it tends to get stuck in the different modes, while the CC algorithm has the highest performance at a computational complexity that is comparable to that of the FCC algorithm in this case (due to MATLAB's very efficient Gaussian random number generator). Qualitatively, similar outcomes are obtained if the parametrisations of the target distribution or the pseudo-priors are changed.

## 5.2 Partially observed mixture variables

In this example we consider a model with two layers, where a pair  $Y = (M, Z)$  of random variables, forming a mixture model  $\tilde{\pi}$  on  $\mathbf{Y} = \llbracket 1, n \rrbracket \times \mathbf{Z}$  of the form described in Section 2.2, is only *partially observed* through some random variable  $X$  taking values in some other state space  $(\mathbf{X}, \mathcal{X})$ . More specifically, we assume that the distribution of  $X$  conditionally on  $Y$  is given by some Markov transition density  $g$  on  $\mathbf{Y} \times \mathbf{X}$ , i.e.

$$X \mid (M, Z) \sim g((M, Z), x) \lambda(dx),$$

where  $\lambda$  is some reference measure on  $(\mathbf{X}, \mathcal{X})$ . When operating on a model of this form one is typically interested in computing the conditional distribution of the latent variable  $Y$  given some distinguished value  $X = x \in \mathbf{X}$  of the observed variable. This posterior distribution has the density

$$\pi^*(m, z \mid x) = \frac{g((m, z), x) \tilde{\pi}(m, z)}{\iint g((m, z), x) \tilde{\pi}(m, z) |dm| \nu(dz)} \quad ((m, z, x) \in \mathbf{Y} \times \mathbf{X})$$

w.r.t. the product  $|dm| \nu(dz)$ . Since the observation  $x$  is fixed, we simply omit this quantity from the notation and write  $\pi^*(m, z \mid x) = \pi^*(m, z)$ . Note that  $\pi^*$  is again a mixture model on  $\mathbf{Y}$ , and our objective is to sample this distribution.

In order to evaluate, in this framework, the performances of the MCMC samplers discussed in the previous section we let, as in the previous example,  $\mathbf{Y} = \llbracket 1, 2 \rrbracket \times \mathbb{R}$  and consider the Gaussian mixture model

$$\tilde{\pi}(m, z) = \alpha_m \phi(z; \mu_m, \sigma^2) \quad ((m, z) \in \llbracket 1, 2 \rrbracket \times \mathbb{R}),$$

where  $\alpha_1 = 1/4$ ,  $\alpha_2 = 3/4$ ,  $\mu_1 = -1$ ,  $\mu_2 = 1$ , and  $\sigma = \sqrt{2}$ . (Note that letting  $\alpha_1 = \alpha_2 = 1/2$  yields the mixture model (13) of the previous example.) In addition, we let be  $(M, Z)$  be partially observed through

$$X = Z^2 + \varsigma \varepsilon, \quad (14)$$

where  $\varsigma = \sqrt{1}$  and  $\varepsilon$  is a standard Gaussian noise variable which is independent of  $Z$ . Consequently, the measurement density (with respect to Lebesgue measure) is given by  $g((m, z), x) = \phi(x; z^2, \varsigma^2)$ ,  $x \in \mathbb{R}$ , in this case. For the fixed observation value  $x = .4$  we estimated the posterior distribution

$$\pi^*(m, z) \propto \alpha_m \phi(z; \mu_m, \sigma^2) \phi(x; z^2, \varsigma^2) \quad ((m, z) \in \llbracket 1, 2 \rrbracket \times \mathbb{R})$$

and the corresponding posterior mean  $\mu_z := \int z \pi^*(dz)$  using Algorithm 3 and Algorithm 4. Note that we are unable to sample directly the conditional distribution  $\pi^*(z \mid m)$  in this case due to the nonlinearity of the observation equation (14); thus, Algorithm 2 is excluded from our comparison. In addition, we implemented the Gibbs sampler Algorithm 1 with Step (ii) replaced by a Metropolis-Hastings operation, yielding a Metropolis-within-Gibbs (MwG) sampler. This Metropolis-Hastings operation as well as in the corresponding operation in Step (iii) of the MCC sampler (Algorithm 3) used the conditional prior distribution as proposal, e.g.

$$R_\ell(u, dz) = \tilde{\pi}(dz \mid m) \quad ((\ell, u) \in \llbracket 1, 2 \rrbracket \times \mathbb{R}).$$

This distribution was also used for designing the pseudo-priors in the MCC and FCC algorithms, e.g.,

$$\rho_\ell(dz) = \tilde{\pi}(dz \mid m) \quad ((\ell, u) \in \llbracket 1, 2 \rrbracket \times \mathbb{R}),$$

and consequently the MCC sampler can, as in the previous example, be viewed as a “random refreshment”-version (using the terminology of [3]) of the FCC sampler. The resulting autocorrelation function estimates are displayed in Figure 2, which shows that the FCC and MCC algorithms are clearly superior, in terms of autocorrelation, to the MwG sampler. Even though the MCC sampler has, as expected from Theorem 5, a small advantage to the FCC sampler in terms of autocorrelation, both samplers exhibit, very similar mixing properties. This is particularly appealing in the light of the



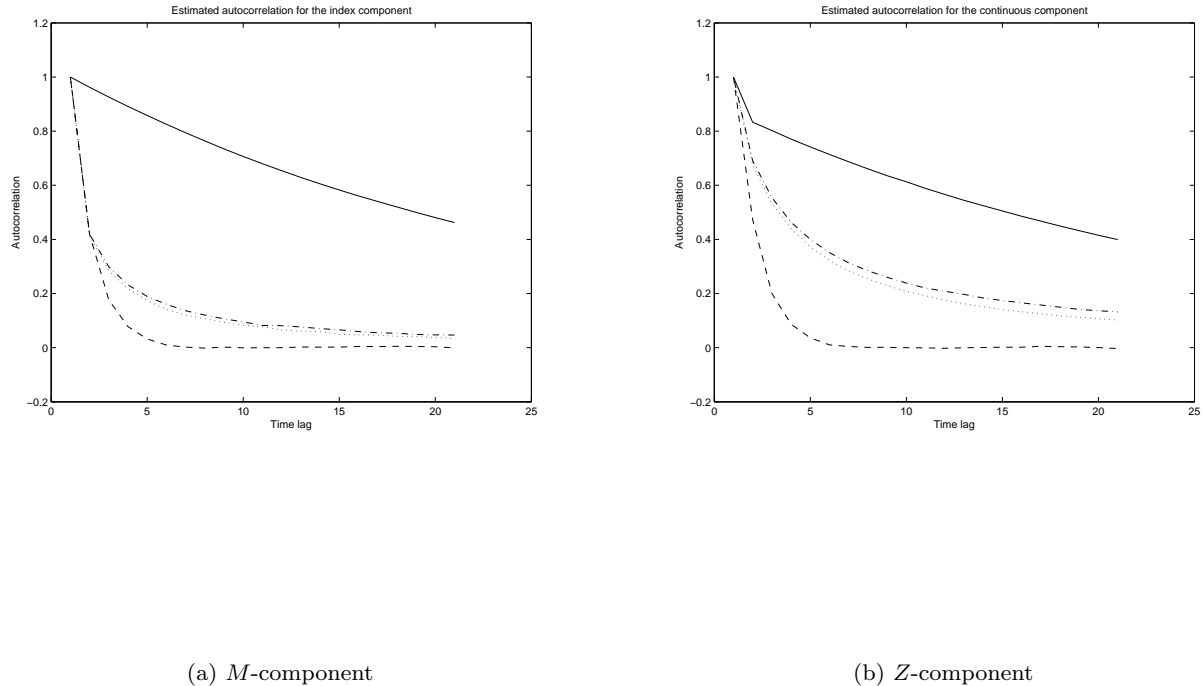


Fig. 1: Plot of estimated autocorrelation for the standard Gibbs sampler (solid line), Algorithm 2 (dashed line), Algorithm 3 (dotted line), and Algorithm 4 (dash-dotted line) when applied to the model (13).

Table 1: Posterior means delivered the MwG, MCC, and FCC algorithms for the partially observed mixture model (14) together with the corresponding CPU times. The true posterior mean (for  $x = .4$ ) is  $\mu_z = .315$ .

algorithm	mean	CPU time (s)
MwG	.334	50.9
MCC	.311	58.7
FCC	.314	33.4

CPU times reported in Table 1, which shows that the FCC sampler almost to twice as fast as the MCC sampler for our implementation. Table 1 reports also the posterior mean estimates obtained with the different algorithms, and apparently the slow mixing of the MwG sampler rubs off on the precision of the corresponding estimate. The true value,  $\mu_z = .315$ , was obtained using numerical integration.

Figure 3 displays the estimate of the marginal posterior density  $\pi^*(z)$  obtained by applying a Gaussian

kernel smoothing function to the output of the FCC algorithm. The exact posterior, obtained using numerical integration, is plotted for comparison.

Finally, we remark that also the results obtained in this example appear to be relatively insensitive to the parametrisation of the model and the pseudo-priors.

## 6 Conclusion

We have compared some data-augmentation-type MCMC algorithms sampling from mixture models comprising a discrete as well as a continuous component. By casting Carlin & Chib’s pseudo-prior into our framework we obtained a sampling scheme that is considerably more efficient than the standard Gibbs sampler, which in general exhibits poor state-space exploration due to strong correlation between the discrete and continuous components (as a result of the highly multimodal nature of the mixture model). In the case where simulation of the continuous component  $Z$  conditionally on  $M$  is infeasible, we used a metropolised version of the algorithm,




Fig. 2: Plot of estimated autocorrelation for the Metropolis-within-Gibbs sampler (solid line), Algorithm 3 (dotted line), and Algorithm 4 (dash-dotted line) when applied to the model (14).

referred to as the MCC sampler, that handled this issue by means of an additional Metropolis-Hastings step in the spirit of the hybrid sampler. In this case our simulations indicate, interestingly, that the loss of mixing caused by simply passing, as in the FCC algorithm, the value of the  $M$ th auxiliary variable, generated by sampling from the pseudo-priors at the beginning of the loop, directly to  $Z$  without any additional refreshment is limited. Thus, we consider the FCC algorithm, which we proved to be  $\pi^*$ -reversible, as strong contender to the MCC sampler in terms of efficiency (variance per unit CPU).

Our theoretical results comparing the MCC and FCC samplers deal exclusively with mixing properties of the restriction of the MCMC output to the discrete component, and the extension of these results to the continuous component is left as an open problem. However, we believe that the discrete component is indeed the quantity of interest, as our simulations indicate that the degree mixing of the discrete component gives a limitation of the degree of mixing of the bivariate chain due to the multimodal nature of the mixture.

There are several possible improvements of the FCC algorithm. For instance, following [5], only a subset of the pseudo-priors (namely those with indices belonging to some neighborhood of the current  $M$ ) could be sampled at each iteration, yielding a very efficient algorithm from a computational point of view. Such an approach could be also used for handling the case of an infinitely large index space (i.e.  $n = \infty$ ).

## References

1. Carlin, B.P., Chib, S.: Bayesian model choice via Markov chain Monte Carlo methods. J. of the Royal Statistical Society, ser. B, Methodology (1995)
2. Hurn, M., Justel, A., Robert, C.P.: Estimating mixtures of regressions. Journal of Computational and Graphical Statistics **12**(1), 55–79 (2003). DOI 10.1198/1061860031329. URL <http://www.tandfonline.com/doi/abs/10.1198/1061860031329>
3. Maire, F., Douc, R., Olsson, J.: Comparison of asymptotic variances of inhomogeneous Markov chains with applications to Markov chain Monte Carlo methods. The Annals of Statistics (2014). To appear.




Fig. 3: Probability density estimate based on the sequence  $\{Z_k^{(4)}; k \in \llbracket 1001, 10^5 \rrbracket\}$  generated by Algorithm 4 (dashed line) for the partially observed mixture model (14) together with the exact posterior density (solid line).

4. Mira, A.: Ordering and improving the performance of Monte Carlo Markov chains. *Statistical Science* pp. 340–350 (2001)
5. Petralias, A., Dellaportas, P.: An mcmc model search algorithm for regression problems. *Journal of Statistical Computation and Simulation* **83**(9), 1722–1740 (2013)
6. Robert, C.P., Casella, G.: *Monte Carlo Statistical Methods*. Springer (2004)
7. Tierney, L.: A note on Metropolis-Hastings kernels for general state spaces. *Annals of Applied Probability* pp. 1–9 (1998)